

UNITED STATES PATENT APPLICATION

of

LEE S. MIGHDOLL,

BRUCE A. LEAK,

STEPHEN G. PERLMAN,

and

PHILLIP Y. GOLDMAN

for

**A METHOD FOR DOWNLOADING A WEB PAGE TO A
CLIENT FOR EFFICIENT DISPLAY ON A TELEVISION SCREEN**

WORKMAN, NYDEGGER & SEELEY
A PROFESSIONAL CORPORATION
ATTORNEYS AT LAW
1000 EAGLE GATE TOWER
60 EAST SOUTH TEMPLE
SALT LAKE CITY, UTAH 84111

UNITED STATES PATENT AND TRADEMARK OFFICE

BACKGROUND OF THE INVENTION

1. Related Applications

This application is a continuation of U.S. Patent Application Serial No. 09/095,457, filed June 10, 1998, entitled, "A Method for Downloading a Web Page to a Client for Efficient Display on a Television Screen," which is hereby incorporated by reference.

2. The Field of the Invention

The present invention pertains to the field of client-server computer networking. More particularly, the present invention relates to a method and apparatus for providing proxying and document transcoding in a server in a computer network.

3. Background and Related Art

The number of people using personal computers has increased substantially in recent years, and along with this increase has come an explosion in the use of the Internet. One particular aspect of the Internet which has gained widespread use is the World-Wide Web ("the Web"). The Web is a collection of formatted hypertext pages located on numerous computers around the world that are logically connected by the Internet. Advances in network technology and software providing user interfaces to the Web ("Web browsers") have made the Web accessible to a large segment of the population. However, despite the growth in the development and use of the Web, many people are still unable to take advantage of this important resource.

Access to the Web has been limited thus far mostly to people who have access to a personal computer. However, many people cannot afford the cost of even a relatively inexpensive personal computer, while others are either unable or unwilling to learn the basic computer skills that are required to access the Web. Furthermore, Web browsers in the prior art

1 generally do not provide the degree of user-friendliness desired by some people, and many
2 computer novices do not have the patience to learn how to use the software. Therefore, it would
3 be desirable to provide an inexpensive means by which a person can access the Web without the
4 use of a personal computer. In particular, it would be desirable for a person to be able to access
5 the Web pages using an ordinary television set and a remote control, so that the person feels
6 more as if he or she is simply changing television channels, rather than utilizing a complex
7 computer network.

8 Prior art Web technology also has other significant limitations which can make a person's
9 experience unpleasant when browsing the Web. Web documents are commonly written in
10 HTML (hypertext Mark-up Language). HTML documents sometimes contain bugs (errors) or
11 have features that are not recognized by certain Web browsers. These bugs or quirks in a
12 document can cause a Web browser to fail. Thus, what is needed is a means for reducing the
13 frequency with which client systems fail due to bugs or quirks in HTML documents.

14 Another problem associated with browsing the Web is latency. People commonly
15 experience long, frustrating delays when browsing the Web. It is not unusual for a person to have
16 to wait minutes after selecting a hypertext link for a Web page to be completely downloaded to
17 his computer and displayed on his computer screen. There are many possible causes for latency,
18 such as heavy communications traffic on the Internet and slow response of remote servers.
19 Latency can also be caused by Web pages including images. One reason for this effect is that,
20 when an HTML document references an image, it takes time to retrieve the image itself after the
21 referencing document has been retrieved. Another reason is that, in the prior art, if the
22 referencing document does not specify the size of the image, the client system generally cannot
23 display the Web page until the image itself has been retrieved. Numerous others sources of
24

1 latency exist with respect to the Web. Therefore, what is needed is a means for reducing such
2 latency, to eliminate some of the frustration which typically has been associated with browsing
3 the Web.

4 Security is another concern associated with the Internet. Internet service providers (ISPs)
5 generally maintain certain information about each customer in a database. This information may
6 include information which a customer may not wish to become publicly known, such as social
7 security numbers and credit card numbers. Maintaining the confidentiality of this information in
8 a system that is connected to an expensive publicly-accessible computer network like the Internet
9 can be problematic. Further, the problem can be aggravated by the fact that an ISP often provides
10 numerous different services, each of which has access to this database. Allowing access to the
11 database by many different entities creates many opportunities for security breaches to occur.
12 Therefore, what is needed is a way to improve the security of confidential customer information
13 in a server system coupled to the Internet.

SUMMARY OF THE INVENTION

An improved method for providing a document to a client coupled to a server is disclosed. In the method, a document is provided to a proxying server. The document includes image data and non-image data that causes the client to generate a display. The document is partitioned into a plurality of partitions. The data of the first partition is downloaded to the client after the data of the first partition is downloaded; the data of the next partition is downloaded to the client. These steps are repeated until each one of the plurality of partitions has been downloaded to the client. Other features of the present invention will be apparent from the accompanying drawings and from the detailed description which follows. Method is described of providing a document to a client coupled to a server.

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention is illustrated by way of example and not limitation in the figures of the accompanying drawings, in which like references indicate similar elements and in which:

Figure 1 illustrates several clients connected to a proxying server in a network.

Figure 2 illustrates a client according to the present invention.

Figure 3 is a block diagram of a server according to the present invention.

Figure 4A illustrates a Server including a proxy cache and a transcoder.

Figure 4B illustrates databases used in a server according to the present invention.

Figure 5 is a flow diagram illustrating a routine for transcoding a document retrieved from a remote server using data stored in a persistent database.

Figure 6 is a flow diagram illustrating a routine for transcoding an HTML document for purposes of eliminating bugs or undesirable features.

Figure 7A is a flow diagram illustrating a routine for reducing latency when downloading a document referencing an image to a client.

Figure 7B is a flow diagram illustrating a routine for efficiently downloading a document to a client for efficient display of a Web page on a television screen.

Figure 8 is a flow diagram illustrating a routine for updating documents stored in the proxy cache using data stored in a persistent database.

Figure 9 is a flow diagram illustrating a routine used by a server for retrieving documents from another remote server.

Figure 10 is a block diagram of a prior art server system showing a relationship between various services and a database.

1 Figure 11 is a block diagram of a server system according to the present invention showing
2 a relationship between various services and a user database.

3 Figure 12 is a flow diagram illustrating a routine used by a server for regulating access to
4 various services provided by the server.
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24

DETAILED DESCRIPTION

A method and apparatus are described for providing proxying and transcoding of documents in a network. In the following description, for purposes of explanation, numerous specific details are set forth in order to provide a thorough understanding of the present invention. It will be evident, however, to one skilled in the art that the present invention may be practiced without these specific details. In other instances, well-known structures and devices are shown in block diagram form in order to avoid unnecessarily obscuring the present invention.

The present invention includes various steps, which will be described below. The steps can be embodied in machine-executable instructions, which can be used to cause a general-purpose or special-purpose processor programmed with the instructions to perform the steps. Alternatively, the steps of the present invention might be performed by specific hardware components that contain hardwired logic for performing the steps, or by any combination of programmed computer components and custom hardware components.

I. System Overview

The present invention is included in a system, known as WebTV™, for providing a user with access to the Internet. A user of a WebTV™ client generally accesses a WebTV™ server via a direct-dial telephone (POTS, for "plain old telephone service"), ISDN (Integrated Services Digital Network), or other similar connection, in order to browse the Web, send and receive electronic mail (e-mail), and use various other WebTV™ network services. The WebTV™ network services are provided by WebTV™ servers using software residing within the WebTV™ servers in conjunction with software residing within a WebTV™ client.

Figure 1 illustrates a basic configuration of the WebTV™ network according to one embodiment. A number of WebTV™ clients 1 are coupled to a modem pool 2 via direct-dial, bi-

directional data connections 29, which may be telephone (POTS, i.e., "plain old telephone service"), ISDN (Integrated Services Digital Network), or any other similar type of connection. The modem pool 2 is coupled typically through a router, Such as that conventionally known in the art, to a number of remote servers 4 via a conventional network infrastructure 3, such as the Internet. The WebTV™ system also includes a WebTV™ server 5, which specifically supports the WebTV™ clients 1. The WebTV™ clients 1 each have a connection to the WebTV™ server 5 either directly or through the modem pool 2 and the Internet 3. Note that the modem pool 2 is a conventional modem pool such as those found today throughout the world providing access to the Internet and private networks.

Note that in this description, in order to facilitate explanation the WebTV™ server 5 is generally discussed as if it were a single device, and functions provided by the WebTV™ services are generally discussed as being performed by such single device. However, the WebTV™ server 5 may actually comprise multiple physical and logical devices connected in a distributed architecture, and the various functions discussed below which are provided by the WebTV™ services may actually be distributed among multiple WebTV™ server devices.

II. Client System

Figure 2 illustrates a WebTV™ client 1. The WebTV™ client 1 includes an electronics unit 10 (hereinafter referred to as "the WebTV™ box 10"), an ordinary television set 12, and a remote control 11. In an alternative embodiment of the present invention, the WebTV™ box 10 is built into the television set 12 as an integral unit. The WebTV™ box 10 includes hardware and software for providing the user with a graphical user interface, by which the user can access the WebTV™ network services, browse the Web, send e-mail, and otherwise access the Internet.

1 The WebTV™ client 1 uses the television set 12 as a display device. The WebTV™ box
2 10 is coupled to the television set 12 by a video link 6. The video link 6 is an RF (radio
3 frequency), S-video, composite video, or other equivalent form of video link. In the preferred
4 embodiment, the client 1 includes both a standard modem and an ISDN modem, such that the
5 communication link 29 between the WebTV™ box 10 and the server 5 can be either a telephone
6 (POTS) connection 29a or an ISDN connection 29b. The WebTV™ box 10 receives power
7 through a power line 7.

8 Remote control 11 is operated by the user in order to control the WebTV™ client 1 in
9 browsing the Web, sending e-mail, and performing other Internet-related functions. The
10 WebTV™ box 10 receives commands from remote control 11 via an infrared (IR)
11 communication link. In alternative embodiments, the link between the remote control 11 and the
12 WebTV™ box 10 may be RF or any equivalent mode of transmission.

13 III. Server System

14 The WebTV™ server 5 generally includes one or more computer systems generally having
15 the architecture illustrated in Figure 3. It should be noted that the illustrated architecture is only
16 exemplary; the present invention is not constrained to this particular architecture. The illustrated
17 architecture includes a central processing unit (CPU) 50, random access memory (RAM) 51,
18 read-only memory (ROM) 52, a mass storage device 53, a modem 54, a network interface card
19 (NIC) 55, and various other input/output (I/O) devices 56. Mass storage device 53 includes a
20 magnetic, optical, or other equivalent storage medium. I/O devices 56 may include any or all of
21 devices such as a display monitor, keyboard, cursor control device, etc... Modem 54 is used to
22 communicate data to and from remote servers 4 via the Internet.

1 As noted above, the WebTV™ server 5 may actually comprise multiple physical and
2 logical devices connected in a distributed architecture. Accordingly, NIC 55 is used to provide
3 data communication with other devices that are part of the WebTV™ services. Modem 54 may
4 also be used to communicate with other devices that are part of the WebTV™ services and
5 which are not located in close geographic proximity to the illustrated device.

6 According to the present invention, the WebTV™ server 5 acts as a proxy in providing the
7 WebTV™ client 1 with access to the Web and other WebTV™ services. More specifically,
8 WebTV™ server 5 functions as a “caching proxy.”

9 Figure 4A illustrates the caching feature of the WebTV™ server 5. In Figure 4A, the
10 WebTV™ server 5 is functionally located between the WebTV™ client 1 and the Internet
11 infrastructure 3. The WebTV™ server 5 includes a proxy cache 65 which is functionally coupled
12 to the WebTV™ client 1. The proxy cache 65 is used for temporary storage of Web documents,
13 images, and other information which is frequently used by either the WebTV™ client 1 or the
14 WebTV™ server 5.

15 A document transcoder 66 is functionally coupled between the proxy cache 65 and the
16 Internet infrastructure 3. The document transcoder 66 includes software which is used to
17 automatically revise the code of Web documents retrieved from the remote servers 4, for
18 purposes which are described below.

19 The WebTV™ service provides a document database 61 and a user database 62, as
20 illustrated in Figure 4B. The user database 62 contains information that is used to control certain
21 features relating to access privileges and capabilities of the user of the client 1. This information
22 is used to regulate initial access to the WebTV™ service, as well as to regulate access to the
23 individual services provided by the WebTV™ system, as will be described below. The document
24

1 database 61 is a persistent database which stores certain diagnostic and historical information
2 about each document and image retrieved by the server 5, as is now described.

3 A. Document Database

4 The basic purpose of the document database 61 is that, after a document has once been
5 retrieved by the server the stored information can be used by the server 5 to speed up processing
6 and downloading of that document in response to all future requests for that document. In
7 addition, the transcoding functions and various other functions of the WebTV™ service are
8 facilitated by making use of the information stored in the document database 61, as will be
9 described below. Referring now to Figure 5, the server 5 initially receives a document request
10 from a client 1 (step 501). The document request will generally result from the user of the client
11 1 activating a hypertext anchor (link) on a Web page.

12 The act of activating a hypertext anchor may consist of clicking on underlined text in a
13 displayed Web page using a mouse, for example. The document request will typically (but not
14 always) include the URL (Uniform Resource Locator) or other address of the selected anchor.
15 Upon receiving the document request, the server 5 optionally accesses the document database 62
16 to retrieve stored information relating to the requested document (step 502). It should be noted
17 that the document database 62 is not necessarily accessed in every case. The information
18 retrieved from the document database 62 is used by the server 5 for determining, among other
19 things, how long a requested document has been cached and/or whether the document is still
20 valid. The criteria for determining validity of the stored document are discussed below. The
21 server 5 retrieves the document from the cache 65 if the stored document is valid; otherwise, the
22 server 5 retrieves the document from the appropriate remote server 4 (step 503). The server 5
23
24

1 automatically transcodes the document as necessary based on the information stored in the
2 document database 61 (step 503). The transcoding functions are discussed further below.

3 The document database 61 includes certain historical and diagnostic information for every
4 Web page that is accessed at any time by a WebTV™ client 1. As is well known, a Web page
5 may correspond to a document written in a language such as HTML (Hypertext Mark-Up
6 Language), VRML (Virtual Reality Modeling Language), or another suitable language.
7 Alternatively, a Web page may represent an image, or a document which references one or more
8 images. According to the present invention, once a document or image is retrieved by the
9 WebTV™ server 5 from a remote server 4 for the first time, detailed information on this
10 document or image is stored permanently in the document database 61. More specifically, for
11 every Web page that is retrieved from a remote server 4, any or all of the following data are
12 stored in the document database 61:

- 13 1) information identifying bugs (errors) or quirks in the Web page, or undesirable
14 effects caused when the Web page is displayed by a client 1;
- 15 2) relevant bug-finding algorithms;
- 16 3) the date and time the Web page was last retrieved;
- 17 4) the date and time the Web page was most recently altered by the author;
- 18 5) a checksum for determining whether the Web page has been altered;
- 19 6) the size of the Web page (in terms of memory);
- 20 7) the type of Web page (e.g., HTML document, image, etc.);
- 21 8) a list of hypertext anchors (links) in the Web page and corresponding
22 URLs;

- 1 9) a list of the most popular anchors based on the number of "hits" (requests from a
2 client 1);
- 3 10) a list of related Web pages which can be prefetched;
- 4 11) whether the Web page has been redirected to another remote server 4;
- 5 12) a redirect address (if appropriate);
- 6 13) whether the redirect (if any) is temporary or permanent, and if permanent, the
7 duration of the redirect;
- 8 14) if the Web page is an image, the size of the image in terms of both physical
9 dimensions and memory space;
- 10 15) the sizes of in-line images (images displayed in text) referenced by the document
11 defining the Web page;
- 12 16) the size of the largest image referenced by the document;
- 13 17) information identifying any image maps in the Web page;
- 14 18) whether to resize any images corresponding to the Web page;
- 15 19) an indication of any forms or tables in the Web page;
- 16 20) any unknown protocols;
- 17 21) any links to "dead" Web pages (i.e., pages which are no longer active);
- 18 22) the latency and throughput of the remote server 4 on which the Web page is
19 located;
- 20 23) the character set of the document;
- 21 24) the vendor of the remote server 4 on which the Web page is located;
- 22 25) the geographic location of the remote server 4 on which the Web page is located;
- 23 26) the number of other Web pages which reference the subject Web page;
- 24

- 27) the compression algorithm used by the image or document;
- 28) the compression algorithm chosen by the transcoder;
- 29) a value indicating the popularity of the Web page based on the number of hits by clients; and
- 30) a value indicating the popularity of other Web pages which reference the subject Web page.

B. Transcoding

As mentioned above, the WebTV™ services provide a transcoder 66, which is used to rewrite certain portions of the code in an HTML document for various purposes. These purposes include: (1) correcting bugs in documents; (2) correcting undesirable effects which occur when a document is displayed by the client 1; (3) improving the efficiency of transmission of documents from the server 5 to the client 1; (4) matching hardware decompression technology within the client 1; (5) resizing images to fit on the television set 12; (6) converting documents into other formats to provide compatibility; (7) reducing latency experienced by a client 1 when displaying a Web page with in-line images (images displayed in text); and, (8) altering documents to fit into smaller memory spaces.

There are three transcoding modes used by the transcoder 66: (1) streaming, (2) buffered, and (3) deferred. Streaming transcoding refers to the transcoding of documents on a line-by-line basis as they are retrieved from a remote server 4 and downloaded to the client 1 (i.e., transcoding "on the fly"). Some documents, however, must first be buffered in the WebTV™ server 5 before transcoding and downloading them to the client 1. A document may need to be buffered before transmitting it to the client 1 if the type of changes to be made can only be made after the entire document has been retrieved from the remote server 4. Because the process of

1 retrieving and downloading a document to the client 1 increases latency and decreases
2 throughput, it is not desirable to buffer all documents. Therefore, the transcoder 66 accesses and
3 uses information in the document database 61 relating to the requested document to first
4 determine whether a requested document must be buffered for purposes of transcoding, before
5 the document is retrieved from the remote server 4.

6 In the deferred mode, transcoding is deferred until after a requested document has been
7 downloaded to a client 1. The deferred mode therefore reduces latency experienced by the client
8 1 in receiving the document. Transcoding may be performed immediately after downloading or
9 any time thereafter. For example, it may be convenient to perform transcoding during periods of
10 low usage of WebTV™ services, such as at night. This mode is useful for certain types of
11 transcoding which are not mandatory.

12 1. Transcoding for Bugs and Quirks

13 One characteristic of some prior art Web browsers is that they may experience failures
14 (“crashes”) because of bugs or unexpected features (“quirks”) that are present in a Web
15 document. Alternatively, quirks in a document may cause an undesirable result, even though the
16 client does not crash. Therefore, the transcoding feature of the present invention provides a
17 means for correcting certain bugs and quirks in a Web document. To be corrected by the
18 transcoder 66, bugs and quirks must be identifiable by software running on the server 5.
19 Consequently, the transcoder 66 will generally only correct conditions which have been
20 previously discovered, such as those discovered during testing or reported by users. Once a bug
21 or quirk is discovered, however, algorithms are added to the transcoder 66 to both detect the bug
22 or quirk in the future in any Web document and to automatically correct it.

1 There are countless possibilities of bugs or quirks which might be encountered in a Web
2 document. Therefore, no attempt will be made herein to provide an exhaustive list. Nonetheless,
3 some examples may be useful at this point. Consider, for example, an HTML document that is
4 downloaded from a remote server 4 and which contains a table having a width specified in the
5 document as "0." This condition might cause a failure if the client were to attempt to display the
6 document as written. This situation therefore, can be detected and corrected by the transcoder 66.
7 Another example is a quirk in the document which causes quotations to be terminated with too
8 many quotation marks. Once the quirk is first detected and an algorithm is written to recognize it,
9 the transcoder 66 can automatically correct the quirk in any document.

10 If a given Web document has previously been retrieved by the server 5, there will be
11 information regarding that document available in the document database 61 as described above.
12 The information regarding this document will include whether or not the document included any
13 bugs or quirks that required transcoding when the document was previously retrieved. The
14 transcoder 66 utilizes this information to determine whether (1) the document is free of bugs and
15 quirks, (2) the document has bugs or quirks which can be remedied by transcoding on the fly, or
16 (3) the document has bugs or quirks which cannot be corrected on the fly (i.e., buffering is
17 required).

18 Figure 6 illustrates a routine for transcoding a Web document for purposes of eliminating
19 bugs and quirks. Initially, the server 5 receives a document request from the client 1 (step 601).
20 Next, the document database 61 is accessed to determine whether or not the requested document
21 has been previously retrieved (step 602). If the document has not been previously retrieved, then
22 the server 5 retrieves the document from the remote server 4 (step 609). Next, the retrieved
23 document is analyzed for the presence of bugs or unusual conditions (step 610). Various
24

1 diagnostic information is then stored in the document database 61 as a result of the analysis to
2 note any bugs or quirks that were found (step 611). If any bugs or quirks were found which can
3 be corrected by the transcoder 66, the document is then transcoded and saved to the proxy cache
4 65 (step 612). The transcoded document is then downloaded to the client 1 (step 613). It should
5 be noted that transcoding can be deferred until after the document has been downloaded, as
6 described above; hence, the sequence of Figure 6 is illustrative only.

7 If (in step 602) the requested document had been previously retrieved, then it is determined
8 whether the requested document is still valid (step 603) and whether the document is present in
9 the proxy cache 65 (step 604). If the document is no longer valid, then the document is retrieved
10 from the remote server 4, analyzed for bugs and quirks, transcoded as required, and then
11 downloaded to the client 1 as described above (steps 610-613, step 607). Methods for
12 determining validity of a document are discussed below. If the document is still valid (step 603)
13 and the document is present in the cache 65, the document is downloaded to the client 1 in its
14 current form (as it is stored in the cache), since it has already been transcoded (step 608).

15 The document, however, may be valid but not present in the cache. This may be the case,
16 for example, if the document has not been requested recently and the cache 65 has become too
17 full to retain the requested document. In that case, the document is retrieved again from the
18 remote server 4 (step 605) and then transcoded on the basis of the previously acquired diagnostic
19 information stored within the database 61 for that document. The document is then saved to the
20 cache 65 (step 606). Note that because the document is still valid, it is assumed that the
21 diagnostic information stored in the document database 61 for that document is still valid and
22 that the transcoding can be performed on the basis of that information. Accordingly, once the
23 document is transcoded, the transcoded document is downloaded to the client 1 (step 607).

Again, note that transcoding can be deferred until after the document has been downloaded in some cases.

The validity of the requested document can be determined based on various different criteria. For example, some, HTML documents specify a date on which the document was created, a length of time for which the document will be valid, or both. The validity determination can be based upon such information. For example, a document, which specifies only the date of creation, can be automatically deemed invalid after a predetermined period of time has passed.

Alternatively, validity can be based upon the popularity of the requested document. "Popularity" can be quantified based upon the number of hits for that document, which is tracked in the document database 61. For example, it might be prudent to simply assign a relatively short period of validity to a document which is very popular and a longer period of validity to a document which is less popular.

Another alternative basis for the validity of a document is the observed rate of change of the document. Again, data in the persistent document database 61 can be used. That is, because the document database 61 stores the date and time on which the document was last observed to change, the server 5 can approximate how often the document actually changes. A document or image which is observed to change frequently (e.g., a weather map or a news page) can be assigned a relatively short period of validity. It will be recognized that numerous other ways of determining validity are possible.

2. Transcoding to Reduce Latency

Another purpose for transcoding is to allow documents requested by a client 1 to be displayed by the client 1 more rapidly. Many HTML documents contain references to "in-line"

1 images, or images that will be displayed in text in a Web page. The normal process used in the
2 prior art to display a Web page having in-line images is that the HTML document referencing the
3 image is first downloaded to the client, followed by the client's requesting the referenced image.
4 The referenced image is then retrieved from the remote server on which it is located and
5 downloaded to the client. One problem associated with the prior art, however, is that the speed
6 with which a complete Web page can be displayed to the user is often limited by the time it takes
7 to retrieve in-line images. One reason for this is that it simply takes time to retrieve the image
8 itself after the referencing document has been retrieved. Another reason is that in the prior art, if
9 the referencing document does not specify the size of the image, the Web page generally cannot
10 be displayed until the image itself has been retrieved. The present invention overcomes these
11 limitations.

12 According to the present invention, information stored in the document database 61
13 regarding the in-line images is used to transcode the referencing document in order to reduce
14 latency in displaying the Web page. Once any document which references an in-line image is
15 initially retrieved by the server 5, the fact that the document references an in-line image is stored
16 in the document database. In addition, the size of the image is determined, either from the
17 document (if specified) or from the image itself and then stored in the document database 61.
18 Consequently, for documents which do not specify the size of their in-line images, the size
19 information stored in the database 61 is then used the next time the document is requested in
20 order to reduce latency in downloading and displaying the Web page.

21 Refer now to Figure 7A, which illustrates a routine for reducing latency when
22 downloading a document referencing an image to a client 1. Assume that a client 1 sends a
23 request to the server 5 for an HTML document containing a reference to an in-line image.
24

1 Assume further that the size of the image is not specified in the document itself. Initially, the
2 server 5 determines whether that document has been previously retrieved (step 701). If not, the
3 standard initial retrieval and transcoding procedure is followed (step 706), as described in
4 connection with Figure 6. If, however, the document has been previously retrieved, then the
5 transcoder 66 accesses the size information stored in the document database 61 for the in-line
6 image (step 702). Based on this size information, the HTML document is transcoded such that,
7 when the Web page is initially displayed by the client 1, the area in which the image belongs is
8 replaced by a blank region enveloping the shape of the image (step 703). Thus, any in-line image
9 referenced by a document is displayed initially as a blank region. Consequently, the client 1 can
10 immediately display the Web page corresponding to the HTML document even before the
11 referenced image has been retrieved or downloaded (i.e., even before the size of the image is
12 known to the client 1).

13 As the transcoded HTML document is downloaded to the client, the image is retrieved
14 from the appropriate remote server 4 (step 704). Once the image is retrieved from the remote
15 server 4 and downloaded to the client 1, the client 1 replaces the blank area in the Web page with
16 the actual image (step 705).

17 3. Transcoding to Display Web Pages on a Television

18 As noted above, the client 1 utilizes an ordinary television set 12 as a display device.
19 However, images in Web pages are generally formatted for display on a computer monitor, not a
20 television set. Consequently, the transcoding function of the present invention is used to resize
21 images for display on the television set 12. This includes rescaling images as necessary to avoid
22 truncation when displayed on the television set 12.

1 It should be noted that prior art Web browsers which operate on computer monitors
2 typically use resizable windows. Hence, the size of the visible region varies from client to client.
3 However, because the web browser used by the WebTV™ client 1 is specifically designed for
4 display on a television set, the present invention allows documents and images to be formatted
5 when they are cached.

6 As mentioned previously, prior art servers generally download all of the HTML data, or
7 non-image data, of a Web page to a client and then subsequently download the image data to the
8 client when displaying a Web page, without consideration of the viewable display area of the
9 screen. A trade-off with this prior art approach is that it optimizes total Web page throughput at
10 the expense of the time required to update the viewable display area of the screen. Although this
11 problem may not be as noticeable on computers with large monitors having large display areas,
12 this problem is more noticeable on displays with relatively smaller viewer areas, such as for
13 example a television set 12 coupled to WebTV™ client 1.

14 In another aspect of the present invention, server 5 addresses this issue by taking into
15 consideration the viewable display area of the screen and thus changing the order in which Web
16 based information is downloaded from server 5 to WebTV™ client 1 for the efficient display of
17 a Web page on television set 12. In this embodiment, server 5 reorders the Web page information
18 such that all of the non-image data and image data that appears within the viewable display area
19 of the screen of television set 12 is downloaded before the non-image data and image data of the
20 Web page that is outside the viewable display area of the screen of television set 12 is
21 downloaded. As a result, the overall time required by WebTV™ client 1 to fully generate and
22 display just the portion of the Web page within the viewable display area of television set 12 is
23 reduced.
24

1 Figure 7B is a flow diagram illustrating the steps performed to efficiently download Web
2 page data from server 5 to a WebTV™ client 1 for efficient display within the viewable display
3 area of a television screen in accordance with the teachings of the present invention. Assume that
4 WebTV™ client 1 sends a request to server 5 for a document that is used to generate a Web
5 page. Assume further that the document contains non-image data, such as for example HTML
6 data, and image data. Assume also that the overall size of the Web page is such that the entire
7 Web page cannot be displayed on a single screen of television set 12.

8 Initially, server 5 receives the document that defines the Web page from remote server 4
9 (step 751). After the document has been received from remote server 4, server 5 lays out the
10 entire Web page defined by the received document using well-known techniques (step 753).
11 After the entire Web page has been laid out in server 5, the Web page may be separated or
12 partitioned into a plurality of viewable portions or partitions, such that each partition corresponds
13 to one screen of Web page information viewable on the television screen. In one embodiment,
14 one television screen of Web page information corresponds to the viewable display height of the
15 television screen.

16 In one embodiment, the page height of the entire Web page defined by the document is
17 assumed to equal PH and the display or screen height of television set 12 coupled to WebTV™
18 client 1 is presumed to equal SH (step 755). A variable H is also initialized to a value of 0 (step
19 755). Next, all of the non-image data that drives the layout on WebTV™ client 1 within the
20 viewable display area on television set 12 is downloaded to WebTV™ client 1. In one
21 embodiment, this non-image data that is transmitted or downloaded includes all of the HTML
22 data that drives the layout on WebTV™ client 1 within the partition of the Web page defined by
23 the region H to H + SH - 1 (step 757). After all the non-image data that drives the layout within
24

1 the viewable display area on the screen of television set 12 has been downloaded, server 5 then
2 downloads to WebTV™ client 1 all of the image data that is displayed within the viewable
3 display area of the screen of television Set 12. The image data downloaded from server 5 to
4 WebTV™ client 1 is defined to be the image data displayed by the client within the partition of
5 the Web page defined by H to H + SH -1 (step 759).

6 After the non-image data and image data have been downloaded as described from server 5
7 to WebTV™ client 1 for a first television screen of Web page information, server 5 then
8 sequentially repeats the steps of downloading the non-image data and the image data for each of
9 the remaining viewable television screens of the Web page information until all of the television
10 screens of the Web page have been downloaded from server 5 to WebTV™ client 1. In one
11 embodiment, each remaining television screen of the Web page information is not downloaded
12 until all of the non-image data and image data of a previous television screen Web page have
13 been fully transmitted or downloaded.

14 Referring back to Figure 7B the remaining partitions, or television screens of the Web page
15 information, are downloaded from server 5 to WebTV™ client 1 according to the following
16 remaining steps. After step 759, H is incremented by SH (step 761). If H is less than PH, then
17 processing is looped back to step 757 (step 763). If H is not less than PH, then all of the
18 partitions, or television screens of the Web page, have been downloaded from server 5 to
19 WebTV™ client 1. By reordering the Web page information as described, each television screen
20 of a Web page information is efficiently downloaded from server 5 to WebTV™ client 1 and
21 may therefore be fully generated and displayed on television set 12 in a reduced amount of time.

4. Transcoding for Transmission Efficiency

Documents retrieved by the server 5 are also transcoded to improve transmission efficiency. In particular, documents can be transcoded in order to reduce high frequency components in order to reduce interlace flicker when they are displayed on a television set. Various methods for coding software or hardware to reduce perceptual interlace flicker are described in U.S. Patent No. 5,862,220.

Documents can also be transcoded in order to lower the resolution of the displayed Web page. Reducing the resolution is desirable, because images formatted for computer Systems will generally have a higher resolution than the NTSC (National Television Standards Committee) video format used by conventional television sets. Since the NTSC video does not have the bandwidth to reproduce the resolution of computer-formatted images, the bandwidth consumed in transmitting images to the client 1 at such a high resolution would be wasted.

5. Other Uses for Transcoding

Transcoding is also used by the present invention to recode a document using new formats into older, compatible formats. Images are often displayed in the JPEG (Joint Picture Experts Group) format or the GIF image format. JPEG often consumes less bandwidth than GIF, however. Consequently, images which are retrieved in GIF format are sometimes transcoded into JPEG format. Methods for generally converting images between GIF and JPEG formats are well known.

Other uses for transcoding include transcoding audio files. For example, audio may be transcoded into different formats in order to achieve a desired balance between memory usage, sound quality, and data transfer rate. In addition, audio may be transcoded from a file format (e.g., an ".AU" file) to a streaming format (e.g., MPEG 1 audio). Yet another use of audio

1 transcoding is the transcoding of MIDI (Musical Instrument Digital Interface) data to streaming
2 variants of MIDI.

3 Additionally, documents or images requiring a large amount of memory (e.g., long lists)
4 can be transcoded in order to consume less memory space in the client 1. This may involve, for
5 example, separating a large document or image into multiple sections. For example, the server 5
6 can insert tags at appropriate locations in the original document so that the document appears to
7 the client 1 as multiple Web pages. Hence, while viewing a given page representing a portion of
8 the original document, the user can view the next page (i.e., the next portion of the original
9 document) by activating a button on the screen as if it were an ordinary hypertext anchor.

10 C. Proxying

11 As noted above, the server 5 functions as a proxy on behalf of the client 1 for purposes of
12 accessing the Web. The document database 61 is used in various ways to facilitate this proxy
13 role, as will now be described.

14 1. Updating Cached Documents

15 It is desirable to store frequently requested HTML documents and images in the proxy
16 cache 65 to further reduce latency in providing Web pages to the client 1. However, because
17 some documents and images change over time, documents in the cache 65 will not be valid
18 indefinitely, as mentioned above. A weather map or a news-related Web page, for example, are
19 likely to be updated quite frequently. Consequently, it is desirable for the server 5 to have the
20 ability to estimate the frequency with which documents change, in order to determine how long a
21 document can safely remain within the proxy cache 65 without being updated.

22 The persistent database 65 is used to store the date and time of the last several fetches of
23 each document and image retrieved from a remote server 4, along with an indication of any
24

1 changes that were detected, if any. A document or image which has been stored in the cache 65
2 is then retrieved on a periodic basis to determine if it has been changed. Change status
3 information indicating whether the document has changed since the previous fetch is then stored
4 in the document database 61. If no changes are detected, then the time interval between fetches
5 of this document is increased. If the document has changed, the time interval is maintained or
6 decreased. As a result, items in the cache 65 which change frequently will be automatically
7 updated at frequent intervals, whereas documents which do not change often will be replaced in
8 the cache less frequently.

9 Figure 8 illustrates a routine for updating documents stored in the proxy cache 65 using
10 data stored in the document database 61. Assume a document X has been stored in the proxy
11 cache 65. Document X remains in the cache 65 until a predetermined update period T_1 expires
12 (step 801). Upon the expiration of the update period T_1 , the document X is again retrieved from
13 the appropriate remote server 4 (step 802). The newly-retrieved document X is then compared to
14 the cached version of document X (step 803). If the document has changed, then the cached
15 version of document X is replaced with the newly-retrieved version of document X (step 806). If
16 not, then the update period T_1 is increased according to a predetermined time increment Dt_1 (step
17 804). In any case, the date and time and the change status of document X is saved to the
18 document database 61 (step 805).

19 Document and Image Prefetching

20 The document database 61 is also used by the server 5 to store prefetching information
21 relating to documents and images. In particular, the database stores, for each document that has
22 been retrieved, a list of images referenced by the document, if any, and their locations.
23 Consequently, the next time a document is requested by a client 1, the images can be
24

1 immediately retrieved by the server 5 (from the cache 65, if available, or from the remote server
2 4), even before the client 1 requests them. This procedure improves the speed with which
3 requested Web pages are downloaded to the client.

4 The document database 61 is also used to facilitate a process referred to as "server-advised
5 client prefetching." Server-advised client prefetching allows the server 5 to inform the client 1 of
6 documents or images which are popular to allow the client 1 to perform the prefetching. In
7 particular, for any given document, a list is maintained in the server 5 of the most popular
8 hypertext anchors in that document (i.e., those which have previously received a large number of
9 hits). When that document is requested by the client 1, the server 5 provides the client 1 with an
10 indication of these popular links.

11 3. Redirects

12 Web pages are sometimes forwarded from the remote server on which they are initially
13 placed to a different location. Under the HTTP (Hypertext Transport Protocol), such forwarding
14 is sometimes referred to as a "redirect." When an HTML document is initially stored on one
15 remote server and then later transferred to another remote server, the first remote server will
16 provide, in response to a request for that document, an indication that the document has been
17 transferred to a new remote server. This indication generally includes a forwarding address
18 ("redirect address"), which is generally a URL.

19 In the prior art, when a computer requesting a Web page receives a redirect, it must then
20 submit a new request to the redirect address. Having to submit a second request and wait for a
21 second response consumes time and increases overall latency. Consequently, the present
22 invention uses the document database 61 to store any redirect address for each document or
23 image. Any time a redirected document is requested, the server 5 automatically accesses the
24

1 redirect address to retrieve the document. The document or image is provided to the client 1
2 based on only a single request from the client 1. The change in location of the redirected
3 document or image remains completely transparent to the client 1.

4 Figure 9 illustrates a routine performed by the server 5 in accessing documents which may
5 have been forwarded to a new remote server. Initially, the server 5 receives a request for a
6 document, which generally includes an address (step 901). The server 5 then accesses the
7 document database 61 to determine whether there is a redirect address for the requested
8 document (step 902). If there is no redirect address, then the server 5 accesses a remote server 4
9 based on the address provided in the document request from the client 1 (step 903). Assuming
10 that the remote server 4 does not respond to the server 5 with a redirect (step 904), the document
11 is retrieved and downloaded to the client 1 by the server 5 (step 907). If, however, a redirect
12 address was stored in the document database 61 (step 902), then the server 5 accesses the
13 requested document according to the redirect address (step 906). Or, if the remote server 4
14 responded with a redirect (step 904), then the server 5 saves the redirect address to the document
15 database 61 (step 905) and accesses the requested document according to the redirect address
16 (step 906).

17 4. Other Proxy Functions

18 The document database 61 also stores information relating to the performance of each
19 remote server 4 from which a document is retrieved. This information includes the latency and
20 throughput of the remote server 4. Such information can be valuable in instances where a remote
21 server 4 has a history of responding slowly. For example, when the document is requested, this
22 knowledge can be used by the server 5 to provide a predefined signal to the client 1. The client 1
23
24

1 can, in response to the signal, indicate to the user that a delay is likely and give the user the
2 option of canceling the request.

3 5. Backoff Mode

4 Although the server 5 generally operates in the proxy mode, it can also enter a "backoff
5 mode" in which the server 5 does not act as a proxy, or the server 5 performs only certain aspects
6 of the normal proxying functions. For example, if the proxy cache 65 is overloaded, then the
7 server 5 can enter a backoff mode in which documents are not cached but are transcoded as
8 required. Alternatively, during times when the server 5 is severely overloaded with network
9 traffic, the server 5 may instruct the client 1 to bypass the server 5 and contact remote servers 4
10 directly for a specified time or until further notice. Or, the server 5 can enter a flexible backoff
11 mode in which the client 1 will be instructed to contact a remote server 4 directly only for certain
12 Web sites for a limited period of time.

13 D. Access to WebTV™ Services

14 The WebTV™ server 5 provides various services to the client 1, such as proxying and
15 electronic mail ("e-mail"). In the prior art, certain difficulties are associated with allowing a
16 client computer access to different services of an Internet service, as will now be explained with
17 reference to Figure 10.

18 Figure 10 illustrates a client-server system according to one prior art embodiment. The
19 server 76 provides various services A, B, and C. The server 76 includes a database 71 for storing
20 information on the user's access privileges to services A, B, and C. The client 75 of the
21 embodiment of Figure 10 accesses any of services A, B, and C by contacting that service
22 directly. The contacted service then accesses the database 71, which stores the access privileges
23 of the client 75, to determine whether the client 75 should be allowed to access that service.
24

1 Hence, each service provided by the server 76 requires direct access to the database 71. This
2 architecture results in a large number of accesses being made to the database 71, which is
3 undesirable. In addition, the fact that each service independently has access to the database raises
4 security concerns. Specifically, it can be difficult to isolate sensitive user information. The
5 present invention overcomes such difficulties using a technique which is now described.

6 1. Tickets Containing Privileges And Capabilities

7 As shown in Figure 11, the server 5 provides a number of services D, E, and F, 77, 78 and
8 80, respectively, and a log-in service 78. The log-in service is used specifically to control initial
9 log-on procedures by a client 1. The log-in service 78 has exclusive access to the user database
10 62 (discussed above with respect to Figure 4B). The log-in service 78 and the user database 62
11 are located within a first security zone 84. Service D is located within a second security zone 86,
12 while services E and F are contained within a third security zone 88. Note that the specific
13 arrangement of security zones 84, 86, and 88 with respect to services D, E, and F is illustrative
14 only.

15 The user database 62 of the present invention stores various information pertaining to
16 each authorized user of a client 1. This information includes account information, a list of the
17 WebTV™ services that are available to the particular user, and certain user preferences. For
18 example, a particular user may not wish his client 1 to be used to access Web pages having adult-
19 oriented subject matter. Consequently, the user would request that his account be filtered to
20 prevent access to such material. This request would then be stored as part of the user data in the
21 user database 62.

22 With regard to user preferences, the hypertext links selected by a given user can be
23 tracked, and those having the largest number can be stored in the user database 62. The list can
24

1 then be provided to the client 1 for use in generating a menu screen of the user's favorite Web
2 sites, to allow the user to directly access those Web sites. The list can also be used by the server
3 5 to analyze the user's interests and to formulate and provide to the user a list of new Web sites
4 which the user is likely to be interested in. The list might be composed by associated key words
5 in Web pages selected by the user with other Web pages.

6 Referring again to Figure 11, in response to a log-on request by a client 1, the log-in
7 service 78 consults the user database 62 to determine if access to the server 5 by this particular
8 client 1 is authorized. Assuming access is authorized, the log-in service 78 retrieves certain user
9 information pertaining to this particular client 1 from the user database 62. The log-in service
10 then generates a "ticket" 82, which is an information packet including the retrieved information.
11 The ticket 82 is then provided to the client 1 which requested access.

12 The ticket 82 includes all information necessary to describe the access privileges of a
13 particular user with respect to all services provided by the server 5. For example, the ticket may
14 include the user name registered to the client 1, the e mail address assigned to client 1, and any
15 filtering requested by the user with respect to viewing Web sites. Each time the user requests
16 access to one of the services D, E, or F, the client 1 submits a copy of the ticket 82 to that
17 service. The requested service can then determine from the copy of the ticket 82 whether access
18 to that service by that client 1 is authorized and, if so, any important information relating to such
19 access.

20 None of the services provided by the server 5, other than the log-in service 78, has access
21 to the user database 62. Hence, any security-sensitive information can be isolated within the user
22 database 62 and the log-in service 78. Such isolation allows the individual services provided by
23 the server 5 to be placed within separate "firewalls" (security regions), illustrated as security
24

1 zones 84, 86, and 88. In addition, this technique greatly reduces the number of accesses required
2 to the user database 62 compared to the prior art embodiment illustrated in Figure 10.

3 2. Redundancy of Services and Load Balancing

4 The present invention also includes certain redundancies in the various services provided
5 by the server 5. In particular, a given service (e.g., e-mail) can be provided by more than one
6 physical or logical device. Each such device is considered a "provider" of that Service. If a given
7 provider is overloaded, or if the client 1 is unable to contact that provider, the client 1 can contact
8 any of the other providers of that service. When the server 5 receives a log-in request from a
9 client 1, in addition to generating the above-described ticket 82, the log-in service 78
10 dynamically generates a list of available WebTV™ services and provides this list to the client 1.

11 The server 5 can update the list of services used by any client 1 to reflect services
12 becoming unavailable or services coming on-line. Also, the list of services provided to each
13 client 1 can be updated by the server 5 based upon changes in the loading of the server 5, in
14 order to optimize traffic on the server 5. In addition, a client's list of services can be updated by
15 services other than the log-in service 78, such that one service can effectively introduce another
16 service to the client 1. For example, the e-mail service may provide a client 1 with the name, port
17 number and IP of its address book service. Thus, one service can effectively, and securely within
18 the same chain of trust, introduce another service to the client 1.

19 This list of services includes the name of each service, a port number for the provider of
20 each service, and an IP (Internet Protocol) for each service. Different providers of the same
21 service are designated by the same name, but different port numbers and/or IPs. Note that in a
22 standard URL, the protocol is normally specified at the beginning of the URL, such as
23 "HTTP://www. . . ." under the HTTP protocol. However, according to the present invention, the
24

normal protocol designation (i.e., "HTTP") in the URL is replaced with the name of the service, since the port number and IP for each service are known to the client 1. Hence, the client 1 can access any of the redundant providers of a given service using the same URL. This procedure effectively adds a level of indirection to all accesses made to any WebTV™ Service and automatically adds redundancy to the proxy service. It should also be noted that separate service names can also refer to the same service.

Assume, for example, that the e-mail service provided by the WebTV™ system is designated by the service name “WTV-mailto.” A client 1 can access any provider of this e-mail service using the same URL. The client 1 merely chooses the appropriate port number and IP number to distinguish between providers. If the client 1 is unable to connect to one e-mail provider, it can simply contact the next one in the list.

Thus, at log-in time, a client 1 is provided with both a ticket containing privileges and capabilities as well as a list of service providers, as illustrated in Figure 12. Initially, the log-in service 78 determines whether the user of client 1 is a valid user (step 1201). If not, log-in is denied (step 1205). If the user is a valid user, then the log-in service 78 gathers user information from the user database 62 and generates a ticket 82 (step 1202). The log-in service 78 also generates the above-described list of services (step 1203). The ticket 82 and the list of services are then downloaded to the client 1 (step 1204).

3. Asynchronous Notification to Clients by Server

Another limitation associated with prior art Internet servers is the inability to provide asynchronous notification information to the client in the absence of a request from the client to do so. It would be desirable, for example, for a server to notify a client on its own initiative when a particular Web page has changed or that a particular service is inaccessible. The server 5 of the

1 present invention provides such capability, and the client 1 is configured to receive and; decode
2 such notifications. For example, the client 1 can receive updates of its listing of service providers
3 from the server 5 at various points in time, as already described.

4 Similarly, if a particular service provider becomes unavailable, that fact will be
5 automatically communicated to the client 1. As another example, if e-mail addressed to the user
6 has been received by the server 5, then the server 5 will send a message to the client 1 indicating
7 this fact. The client 1 will then notify the user that e-mail is waiting by a message displayed on
8 the television set 12 or by an LED (light emitting diode) built into the housing of WebTV™ box
9 10.

10 Thus, a method and apparatus have been described for providing proxying and transcoding
11 of documents in a network. Although the present invention has been described with reference to
12 specific exemplary embodiments, it will be evident that various modifications and changes may
13 be made to these embodiments without departing from the broader spirit and scope of the
14 invention as set forth in the claims. Accordingly, the specification and drawings are to be
15 regarded in an illustrative rather than a restrictive sense.

16 What is claimed is:
17
18
19
20
21
22
23
24